



Enabling Grids for E-science

Release updates cycles: a site perspective

A. Forti, D. Cesini, J. Templon
Operations Workshop
Stockholm, 14 June 2007

www.eu-egee.org



- <http://savannah.cern.ch/bugs/?23232>
- <http://savannah.cern.ch/bugs/?25019>
- <http://savannah.cern.ch/bugs/?26033>
- <http://savannah.cern.ch/bugs/?25932>
-
- https://gus.fzk.de/ws/ticket_info.php?ticket=22323
- https://gus.fzk.de/ws/ticket_info.php?ticket=21920
- https://gus.fzk.de/ws/ticket_info.php?ticket=22090
- https://gus.fzk.de/ws/ticket_info.php?ticket=22426
- https://gus.fzk.de/ws/ticket_info.php?ticket=22444
- https://gus.fzk.de/ws/ticket_info.php?ticket=22445
- https://gus.fzk.de/ws/ticket_info.php?ticket=21234
- https://gus.fzk.de/ws/ticket_info.php?ticket=18383
- https://gus.fzk.de/ws/ticket_info.php?ticket=22365
-
- https://gus.fzk.de/pages/ticket_details.php?ticket=21392
- https://gus.fzk.de/pages/ticket_details.php?ticket=21234
- https://gus.fzk.de/pages/ticket_details.php?ticket=21155
- https://gus.fzk.de/pages/ticket_details.php?ticket=21103
- https://gus.fzk.de/pages/ticket_details.php?ticket=21044
- https://gus.fzk.de/pages/ticket_details.php?ticket=20942
- https://gus.fzk.de/pages/ticket_details.php?ticket=20941
- https://gus.fzk.de/pages/ticket_details.php?ticket=19849
- https://gus.fzk.de/pages/ticket_details.php?ticket=19464
- https://gus.fzk.de/pages/ticket_details.php?ticket=18688
- https://gus.fzk.de/pages/ticket_details.php?ticket=18687
- https://gus.fzk.de/pages/ticket_details.php?ticket=17781
- https://gus.fzk.de/pages/ticket_details.php?ticket=16750
- https://gus.fzk.de/pages/ticket_details.php?ticket=16551
- https://gus.fzk.de/pages/ticket_details.php?ticket=16065
- https://gus.fzk.de/pages/ticket_details.php?ticket=11804
-
- EGEE BROADCAST (05/16/2007 01:03 PM) - Bug found in LFC/DPM 1.6.4-3, distributed in gLite Update 24

- Introduction
- Some examples of problems
- Update 24
- Documentation
- Patch journey to PS
- PPS

- **The last 2 or 3 updates have created quite few problems to sites. Our impression is that these problems can be categorized in the following groups:**
 - The updates aren't sufficiently tested before arriving to PS
 - The updates are too often
 - PPS timelines have now been changed and this is seen as a really positive step
 - The documentation between one update and the other doesn't take into account problems occurred in the previous update so it is not possible to install an update without incurring in the problems all the other updates might have had.
 - Components with a different urgency and difficulty of installation are bundled together
 - Recipes and manual fixes are sent to LCG-ROLLOUT or broadcasted without ending up in the documentation which makes life difficult also for more prudent sites.
- **We have asked site and ROC managers to provide more specific examples and possible solutions with the purpose of discussing them with SA3 and SA1 (PPS) and draw some conclusions on how to evolve to present to the TCG.**

- **Simple issues which are discovered for a given production release, and are "easy" to solve, take weeks to be solved.**

- **Example:**
 - GGUS ticket 22365 (23 May 2007, still unanswered). The issue is present in yaim 3.0.1-15, and is still present in the version 3.0.1-17 (release 31 in PPS), so why does it take more than 30 seconds to go to yaim function config_gip and change the line 748 (in the last yaim version 3.0.1-17) ??

 - GlueInformationServiceURL: ldap://hostname -f :2135/mds-vo-name=local,o=grid

 - GlueInformationServiceURL: ldap://hostname -f :2170/mds-vo-name=resource,o=grid

- **Apparently small change.....**
 - On the UI and RBs it is necessary to change the DN of lcg-voms.cern.ch in the configuration
 - Affected VOs are dteam, atlas, cms, alice, lhcb and ops.
- **Deployment Issues:**
 - No notification in the release documentation about the necessity of making this configuration change.
 - The central LFC for dteam (at least) does not recognize proxies signed by lcg-voms.cern.ch.
 - The VOMS DNs given by the yaimtool are wrong
 - <http://lcg-sft.cern.ch/yaimtool/yaimtool.py> ☹☹☹
 - The VOMS DNs given in various YAIM example files are wrong.
 - Final Broadcast was sent too late and the recipe is not yet in the release notes.

- RPMs are released into production with simple bugs that should have been noticed in PPS.
- **Example: the release of DPM 1.6.4 has a broken YAIM configuration (affects node BDII, info system plugin) and uses the wrong path for lcgdm-mapfile.**
 - Almost 3 weeks later the issue is still there.
 - Patch cronology:
 - Released into PPS on the 8 May (Tuesday)
 - Deployed by Friday 11 May by PPS site admins
 - Released into PS the next Monday 14 May without adequate testing due to
 - *Weekend*
 - *PPS was in "sleep mode" (site admins asked to stop the patch flow).*
- **FIX: A more clearly defined set of systematic tests in PPS or in SA3. (Is there any systematic testing framework?)**

- **Part (and only part) of the update was high priority (lcvoms cert update).**
- **In the same update, there were many other things, in particular DPM 1.6.4 which is not something you can do on the fly, if you don't want to disrupt your SE service.**
 - You need to go in downtime even for a short while (i.e. 15 min)
- **On top of it the DPM version provided by update 24 was buggy.**
 - Because people had to install the high priority part, they also had to install the buggy RPM (with a manual workaround). (see later slides)
- **If DPM update was considered high priority because of the planned changes in SAM, it should have been stated clearly.**

- **FIXES:**

- High priority update should contain only the components that are impacted by the high priority update.
- Critical services update with non trivial procedures involving a temporary (even short) shutdown of the service should be released as separate updates too (e.g. DPM release involving a database schema upgrade)
- Non critical updates, in particular updates impacting only the client side of services, can be bundled together in one update.

- **Comment:**

- Continuous release process is positive because it allows flexibility as long as it is clearly stated when updates require planning.

- Documentation in <http://glite.web.cern.ch/glite/documentation> the following three are rather old documents, does it still make sense?
 - Introduction to the Manual Upgrade Procedure
 - gLite v3.0 Advanced Installation and Configuration Guide
 - generic gLite v3.0 VO Configuration Guide (still points to a glite 1.5 guide!!)

- **Just put yourselves in the "skin" of a new site admin!!!!**
 - If he goes to the bottom of that page he finds: How to test your site <http://goc.grid.sinica.edu.tw/gocwiki/HowToTestYourSite> contains old information which should not be used anymore (grid-proxy-init, SFT's, others??)

- **New site admins don't know, or don't pay too much attention to LCG-ROLLOUT, where all the little tidbits are explained and solved.**
 - The issues reported there do not normally get reported in the official documentation.

- **The full set of services which should be running for each node type is not documented.**
 - What is and what is not relevant, configuration files for each daemon and what/where are their log files.
- **Major changes to, e.g., YAIM are not sufficiently highlighted in the release. Finding a major restructuring in a change from 3.0.0->3.0.1 is just not expected.**
 - It would have helped sites a lot to have a worked example of migrating an old style VO configuration in site-info.def to a new style vo.d directory.
 - It would also help to have more sensible version numbers 3.0.1 should have been 4.0.0 with such a structural change.
- **Release notes shouldn't change under sys admins nose!!!**
 - When release notes are upgraded with workarounds, please issue a broadcast.
 - Even better the release notes should contain a "ChangeLog" and all known issues of the current and previous releases. If a given issue is solved in the meantime, only the Changelog will mention it.

- **Releases should NOT contain updates that concern only CERN (even if temporarily) see FTS 2.0 in update 24 whose release notes contain the following sentence:**
 - The current WLCG planning foresees running version FTS 2.0 only at CERN-PROD for few weeks. Tier-1 sites should remain at version 1.5.....
- **Some new "features" don't appear in the release notes at all, such as the R-GMA SAME job wrapper.**
 - Everything should be documented - even if you think it will be invisible to sites.
- **Release notes often do not mention major changes that will affect services in production**
 - Example: introduction of new pools of accounts for prd and sgm users was mentioned only at the bottom of YAIM-3.0.1 guide, not at all in release notes of the update that introduced new YAIM release (see GGUS 20942).

- **Changes are in units of "patches":**
 - <https://savannah.cern.ch/patch/?group=jra1mdw>
- **Those normally are pretty self-contained, although there is some tendency for developers to put more than one change in a patch as it makes the process faster.**
- **The patches all progress through certification and PPS independently, and what gets seen as an "update" in the PS is just whatever bundle of patches happens to make it out of the PPS stage that week.**
- **The patch structure seems to vanish in that final step, presumably in the cause of simplicity, but even looking at the detailed docs for the updates the patch information seems not to be there.**
 - In PPS the patch level information is still preserved.
 - Perhaps we need to request that patch-level information is preserved also in PS so that sites can choose to install individual patches if they want.
- **Another procedural point is that if you look at the patches you can see that they get closed once they go to production - the assumption is that they must be working to have got that far!**
 - Probably need some extra piece of the release process which allows for the possibility that patches may need to be pushed back even after they get to production.

- **Testing and certification seems not to be very thorough, since some very apparent problems appear too often.**
 - There is not enough time for patch testing in PPS due to the weekly updates.
- **The following changes have already been implemented**
 - It has already been decided to increase the time between an update and another from 1 to 2.5 weeks.
 - Since PPS update 31 a special group of sites has been created to pre-test the installation and configuration part of the release to avoid the most macroscopic bugs → PPS pre-deployment

- **Even if the middleware stayed in PPS for a year would anybody notice what's there?**
 - Either very big bugs
 - Or an hypothetical “almost production” quality
- **Insufficient engagement of VOs in PPS mean things can slip into production which are subsequently found to be problematic, e.g., VOViews and sgm/prd pool accounts being problematic for ATLAS production.**
- **FIX: While it's easy to wish away other people's time, experiments probably need to have an engagement strategy with PPS. If they really don't have time to engage with it, then it loses much of its meaning, and SA1 would probably be better redeploying its PPS effort into SA3.**

- **Is PPS supposed to be an "almost production" quality service? Where "Almost production" quality means:**
 - Installation/Upgrade is a smooth task.
 - All daemons/services which were updated are at most restarted (i.e. do not need manual intervention).
 - Re-configuration (YAIM!) runs smoothly, and all daemons/services are successfully restarted.
 - All daemons/services run successfully.
 - Sometimes the services hang, or do not behave properly (they die, etc.).

- **If YES, then it makes sense to monitor all the services through SAM, and have CIC/GGUS tickets on sites failing.**
- **If NO, then it makes no sense to have CIC/GGUS tickets, SAM tests are still quite/very useful.**
 - Give time for PPS to fix the services which fail (to install/upgrade, to configure, to run).
 - We open GGUS tickets for the developers and integrators.
 - I would consider high priority those tickets coming from PPS.

- **Clearly and unfortunately, we have been in the second situation, for a long time now.**

- **Because PPS sites have to be run as production sites, sysadmins are probably less daring in their testing than they would otherwise be.**
 - Now that I'm all green you want me to change the configuration? No way!!!

- **Many Thanks for their valuable comments to:**
 - Graeme Stewart
 - Yves Coppens
 - Barry Macevoy
 - Stephen Burke
 - Antun Balaz
 - Mario David
 - Cristina Aftimiei
 - Simone Dalla Fina
 - Michel Jouvin

???